Search: © The ACM Digital Library O The Guide

+cas +latency+dimm



Nothing Found

Your search for +cas +latency+dimm did not return any results.

You may want to try an Advanced Search for additional options.

Please review the Quick Tips below or for more information see the Search Tips.

Quick Tips

• Enter your search terms in <u>lower case</u> with a space between the terms.

sales offices

You can also enter a full question or concept in plain language.

Where are the sales offices?

• Capitalize proper nouns to search for specific people, places, or products.

John Colter, Netscape Navigator

• Enclose a phrase in double quotes to search for that exact phrase.

"museum of natural history" "museum of modern art"

Narrow your searches by using a + if a search term must appear on a page.

museum +art

• Exclude pages by using a - if a search term <u>must not appear</u> on a page.

museum -Paris

Combine these techniques to create a specific search query. The better your description of the information you want, the more relevant your results will be.

museum +"natural history" dinosaur -Chicago

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Search:

The ACM Digital Library
The Guide

+cas +latency+module



Nothing Found

Your search for +cas +latency+module did not return any results.

You may want to try an Advanced Search for additional options.

Please review the Quick Tips below or for more information see the Search Tips.

Quick Tips

• Enter your search terms in lower case with a space between the terms.

sales offices

You can also enter a full question or concept in plain language.

Where are the sales offices?

 Capitalize <u>proper nouns</u> to search for specific people, places, or products.

John Colter, Netscape Navigator

• Enclose a phrase in double quotes to search for that exact phrase.

"museum of natural history" "museum of modern art"

• Narrow your searches by using a + if a search term <u>must appear</u> on a page.

museum +art

• Exclude pages by using a - if a search term must not appear on a page.

museum -Paris

Combine these techniques to create a specific search query. The better your description of the information you want, the more relevant your results will be.

museum +"natural history" dinosaur -Chicago

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

Terms of Usage Privacy Policy Code of Ethics Contact Us

IEEE HOME ! SEARCH IEEE ! SHOP ! WEB ACCOUNT ! CONTACT IEEE



Membership Public	ations/Services Standards Conferences Careers/Jobs
JEES?	Welcome United States Patent and Trademark Office
Help FAQ Terms IEE	E Peer Review Quick Links Se
Welcome to IEEE Xplore® O- Home O- What Gan I Access? O- Log-out	Your search matched 0 of 0 documents. A maximum of 500 results are displayed, 15 to a page, sorted by Relevance Descending order. Refine This Search:
Tables of Conlents	You may refine your search by editing the current search expression or enterinew one in the text box.
O- Journals & Magazines O- Conference Proceedings	((cas latency +dimm) and (pyr >= 1950 and pyr <= 2 Search Check to search within this result set
O- Standards	Results Key: JNL = Journal or Magazine CNF = Conference STD = Standard
Search - By Author - Basic - Advanced - CrossRef	Results: No documents matched your query.
Member Services	
O- Join IEEE O- Establish IEEE Web Account O- Access the IEEE Member Digital Library	
O- Access the IEEE Enterprise File Cabinet	

Copyright © 2004 IEEE — All rights reserved

Home | Log-out | Journals | Conference Proceedings | Standards | Search by Author | Basic Search | Advanced Search | Join IEEE | Web Account |

New this week | OPAC Linking Information | Your Feedback | Technical Support | Email Alerting | No Robots Please | Release Notes | IEEE Online

Publications | Help | FAQ| Terms | Back to Top

IEEE HOME | SEARCH IEEE | SHOP | WEB ACCOUNT | CONTACT IEEE

Membership Publications/Services Standards Conferences Careers/Jobs



	Welcome United States Patent and Trademark Office	1 1
Help FAQ Terms IEE	EE Peer Review Quick Links ** S	Se
Welcome to IEEE Xplore* - Home - What Can I Access? - Log-out Tables of Contents - Journals & Magazines - Conference Proceedings - Standards	Your search matched 0 of 0 documents. A maximum of 500 results are displayed, 15 to a page, sorted by Relevance Descending order. Refine This Search: You may refine your search by editing the current search expression or entenew one in the text box. [(((cas latency < and > module) and (pyr >= 1950 and Search) Check to search within this result set Results Key:	
Search O- By Author O- Basic O- Advanced O- CrossRef	JNL = Journal or Magazine CNF = Conference STD = Standard Results: No documents matched your query.	
Member Services Join IEEE Establish IEEE Web Account Access the IEEE Member Digital Library		
O- Access the IEEE Enterprise File Cabinet		

Print Format

Home | Log-out | Journals | Conference Proceedings | Standards | Search by Author | Basic Search | Advanced Search | Join IEEE | Web Account |
New this week | OPAC Linking Information | Your Feedback | Technical Support | Email Alerting | No Robots Please | Release Notes | IEEE Online Publications | Help | FAQ| Terms | Back to Top

Copyright © 2004 IEEE — All rights reserved

Search: • The ACM Digital Library • The Guide

+title:cas +title:latency



Nothing Found

Your search for +title:cas +title:latency did not return any results.

You may want to try an Advanced Search for additional options.

Please review the Quick Tips below or for more information see the Search Tips.

Quick Tips

• Enter your search terms in lower case with a space between the terms.

sales offices

You can also enter a full question or concept in plain language.

Where are the sales offices?

 Capitalize <u>proper nouns</u> to search for specific people, places, or products.

John Colter, Netscape Navigator

• Enclose a phrase in double quotes to search for that exact phrase.

"museum of natural history" "museum of modern art"

Narrow your searches by using a + if a search term <u>must appear</u> on a page.

museum +art

• Exclude pages by using a - if a search term <u>must not appear</u> on a page.

museum -Paris

Combine these techniques to create a specific search query. The better your description of the information you want, the more relevant your results will be.

museum +"natural history" dinosaur -Chicago

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

<u>Terms of Usage Privacy Policy Code of Ethics Contact Us</u>



Search: • The ACM Digital Library • The Guide

+abstract:cas +abstract:latency



Nothing Found

Your search for +abstract:cas +abstract:latency did not return any results.

You may want to try an Advanced Search for additional options.

Please review the Quick Tips below or for more information see the Search Tips.

Quick Tips

• Enter your search terms in lower case with a space between the terms.

sales offices

You can also enter a full question or concept in plain language.

Where are the sales offices?

 Capitalize <u>proper nouns</u> to search for specific people, places, or products.

John Colter, Netscape Navigator

• Enclose a phrase in double quotes to search for that exact phrase.

"museum of natural history" "museum of modern art"

Narrow your searches by using a + if a search term <u>must appear</u> on a page.

museum +art

• Exclude pages by using a - if a search term <u>must not appear</u> on a page.

museum -Paris

Combine these techniques to create a specific search query. The better your description of the information you want, the more relevant your results will be.

museum +"natural history" dinosaur -Chicago

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

Terms of Usage Privacy Policy Code of Ethics Contact Us



Search: • The ACM Digital Library • The Guide

+cas +latency

1331131

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used cas latency

Found 474 of 145,831

Sort results

by

Display results

relevance

expanded form

Save results to a Binder ? Search Tips Open results in a new

Try an Advanced Search Try this search in The ACM Guide

window

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10

next Relevance scale ...

Best 200 shown

Architecture: The architecture of the DIVA processing-in-memory chip Jeff Draper, Jacqueline Chame, Mary Hall, Craig Steele, Tim Barrett, Jeff LaCoss, John Granacki, Jaewook Shin, Chun Chen, Chang Woo Kang, Ihn Kim, Gokhan Daglikoca June 2002 Proceedings of the 16th international conference on Supercomputing

Full text available: pdf(295.98 KB)

Additional Information: full citation, abstract, citings, index terms

The DIVA (Data IntensiVe Architecture) system incorporates a collection of Processing-In-Memory (PIM) chips as smart-memory co-processors to a conventional microprocessor. We have recently fabricated prototype DIVA PIMs. These chips represent the first smartmemory devices designed to support virtual addressing and capable of executing multiple threads of control. In this paper, we describe the prototype PIM architecture. We emphasize three unique features of DIVA PIMs, namely, the memory interf ...

Keywords: architecture, memory bandwidth, processing-in-memory

Separating data and control transfer in distributed operating systems Chandramohan A. Thekkath, Henry M. Levy, Edward D. Lazowska

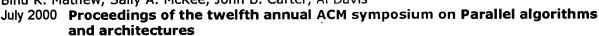
November 1994 Proceedings of the sixth international conference on Architectural support for programming languages and operating systems, Volume 29, 28 Issue 11, 5

Full text available: pdf(1.42 MB)

Additional Information: full citation, abstract, references, citings, index terms

Advances in processor architecture and technology have resulted in workstations in the 100+ MIPS range. As well, newer local-area networks such as ATM promise a ten- to hundred-fold increase in throughput, much reduced latency, greater scalability, and greatly increased reliability, when compared to current LANs such as Ethernet. We believe that these new network and processor technologies will permit tighter coupling of distributed systems at the hardware level, and that distribu ...

Algorithmic foundations for a parallel vector access memory system Binu K. Mathew, Sally A. McKee, John B. Carter, Al Davis



Full text available: pdf(221.23 KB)

Additional Information: full citation, abstract, references, citings, index terms

This paper presents mathematical foundations for the design of a memory controller

subcomponent that helps to bridge the processor/memory performance gap for applications with strided access patterns. The Parallel Vector Access (PVA) unit exploits the regularity of vectors or streams to access them efficiently in parallel on a multi-bank SDRAM memory system. The PVA unit performs scatter/gather operations so that only the elements accessed by the application are tra ...

Constraint analysis for DSP code generation

Bart Mesman, Marino T. J. Strik, Adwin H. Timmer, Jef L. van Meerbergen, Jochen A. G. Jess September 1997 Proceedings of the 10th international symposium on System synthesis

Full text available: pdf(966.00 KB) Publisher Site

Additional Information: full citation, abstract, references, citings

Code generation methods for DSP applications are hampered by the combination of tight timing constraints imposed by the performance requirements of DSP algorithms, and resource constraints imposed by a hardware architecture. In this paper, we present a method to analyze resource- and timing constraints in a single model. The analysis identifies sequencing constraints between operations additional to the precedence constraints. Without the explicit modeling of these sequencing constraints, a sche ...

Measuring Experimental Error in Microprocessor Simulation

Rajagopalan Desikan, Doug Burger, Stephen W. Keckler

June 2001 Proceedings of the 28th annual international symposium on Computer architecture

Full text available: pdf(237.69 KB) Additional Information: full citation, abstract, citings, index terms

Abstract: We measure the experimental error that arises from the use of non-validated simulators in computer architecture research, with the goal of increasing the rigor of simulation- based studies. We describe the methodology that we used to validate a microprocessor simulator against a Compag DS-10L workstation, which contains an Alpha 21264 processor. Our evaluation suite consists of a set of 21 microbenchmarks that stress different aspects of the 21264 microarchitecture. Using the microbenc ...

Measuring experimental error in microprocessor simulation

Rajagopalan Desikan, Doug Burger, Stephen W. Keckler

May 2001 ACM SIGSOFT Software Engineering Notes, Proceedings of the 2001 symposium on Software reusability: putting software reuse in context, Volume 26 Issue 3

Full text available: 📆 pdf(1.03 MB)

Additional Information: full citation, references, index terms

7 An extended classification of inter-instruction dependency and its application in automatic synthesis of pipelined processors

Ing-Jer Huang, Alvin M. Despain

December 1993 Proceedings of the 26th annual international symposium on **Microarchitecture**

Full text available: pdf(1.43 MB)

Additional Information: full citation, references, citings

Keywords: compiler back-end generation, hardware/software tradeoffs, high level synthesis, inter-instruction dependency, pipeline hazard resolution

Scalable lock-free dynamic memory allocation Maged M. Michael

June 2004 ACM SIGPLAN Notices, Proceedings of the ACM SIGPLAN 2004 conference on Programming language design and implementation, Volume 39 Issue 6

Full text available: pdf(213.94 KB) Additional Information: full citation, abstract, references, index terms

Dynamic memory allocators (malloc/free) rely on mutual exclusion locks for protecting the consistency of their shared data structures under multithreading. The use of locking has many disadvantages with respect to performance, availability, robustness, and programming flexibility. A lock-free memory allocator guarantees progress regardless of whether some threads are delayed or even killed and regardless of scheduling policies. This paper presents a completely lock-free memory allocator. It uses ...

Keywords: async-signal-safe, availability, lock-free, malloc

Network behavior: The effectiveness of request redirection on CDN robustness Limin Wang, Vivek Pai, Larry Peterson

December 2002 ACM SIGOPS Operating Systems Review, Volume 36 Issue SI

Additional Information: full citation, abstract, references, citings Full text available: pdf(1.86 MB)

It is becoming increasingly common to construct network services using redundant resources geographically distributed across the Internet. Content Distribution Networks are a prime example. Such systems distribute client requests to an appropriate server based on a variety of factors---e.g., server load, network proximity, cache locality--in an effort to reduce response time and increase the system capacity under load. This paper explores the design space of strategies employed to redirect reque ...

10 A dynamic-SDRAM-mode-control scheme for low-power systems with a 32-bit RISC CPU

Seiji Miura, Kazushige Ayukawa, Takao Watanabe

August 2001 Proceedings of the 2001 international symposium on Low power electronics and design

Full text available: 📆 pdf(955.52 KB) Additional Information: full citation, references, citings, index terms

Keywords: SDRAM controller, active-standby mode, standby mode

11 Hardware-only stream prefetching and dynamic access ordering Chengqiang Zhang, Sally A. McKee

May 2000 Proceedings of the 14th international conference on Supercomputing

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.06 MB) terms

Memory system bottlenecks limit performance for many applications, and computations with strided access patterns are among the hardest hit. The streams used in such applications have extremely poor cache behavior. These access patterns have the advantage of being predictable, though, and this can be exploited to improve the efficiency of the memory subsystem in two ways: memory latencies can be masked by prefetching stream data, and the latencies can be reduced by reordering stream accesses ...

12 A low-cost memory architecture for PCI-based interactive ray casting

Michael Doggett, Michael Meißner, Urs Kanus

July 1999 Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware

Full text available: pdf(1.12 MB) Additional Information: full citation, references, citings, index terms

Keywords: graphics hardware, memory architecture, raycasting, volume rendering accelerator, volume visualization

13 DNS and naming: The design and implementation of a next generation name service for the internet



Venugopalan Ramasubramanian, Emin Gün Sirer

August 2004 Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications

Full text available: pdf(472.93 KB) Additional Information: full citation, abstract, references, index terms

Name services are critical for mapping logical resource names to physical resources in largescale distributed systems. The Domain Name System (DNS) used on the Internet, however, is slow, vulnerable to denial of service attacks, and does not support fast updates. These problems stem fundamentally from the structure of the legacy DNS. This paper describes the design and implementation of the Cooperative Domain Name System (CoDoNS), a novel name service, which provides high lookup performance thro ...

Keywords: DNS, peer to peer, proactive caching

14 Efficient use of memory bandwidth to improve network processor throughput Jahangir Hasan, Satish Chandra, T. N. Vijaykumar

May 2003 ACM SIGARCH Computer Architecture News, Proceedings of the 30th annual international symposium on Computer architecture, Volume 31 Issue 2

Full text available: pdf(184.83 KB) Additional Information: full citation, abstract, references

We consider the efficiency of packet buffers used in packet switches built using network processors (NPs). Packet buffers are typically implemented using DRAM, which provides plentiful buffering at a reasonable cost. The problem we address is that a typical NP workload may be unable to utilize the peak DRAM bandwidth. Since the bandwidth of the packet buffer is often the bottleneck in the performance of a shared-memory packet switch, inefficient use of available DRAM bandwidth further reduces th ...

15 DCAS-based concurrent deques

Ole Agesen, David L. Detlefs, Christine H. Flood, Alexander T. Garthwaite, Paul A. Martin, Nir N. Shavit, Guy L. Steele

July 2000 Proceedings of the twelfth annual ACM symposium on Parallel algorithms and architectures

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(298.15 KB) terms

The computer industry is currently examining the use of strong synchronization operations such as double compare-and-swap (DCAS) as a means of supporting non-blocking synchronization on tomorrow's multiprocessor machines. However, before such a strong primitive will be incorporated into hardware design, its utility needs to be proven by developing a body of effective non-blocking data structures using DCAS. As part of this effort, we present two new linearizable non-blocking impl ...

16 Decentralized storage systems: Taming aggressive replication in the Pangaea widearea file system



Yasushi Saito, Christos Karamanolis, Magnus Karlsson, Mallik Mahalingam December 2002 ACM SIGOPS Operating Systems Review, Volume 36 Issue SI

Additional Information: full citation, abstract, references Full text available: pdf(1.93 MB)

Pangaea is a wide-area file system that supports data sharing among a community of widely distributed users. It is built on a symmetrically decentralized infrastructure that consists of commodity computers provided by the end users. Computers act autonomously to serve data to their local users. When possible, they exchange data with nearby peers to improve the system's overall performance, availability, and network economy. This approach is realized by aggressively creating a replica of a file w ...

17 Relax: A new circuit for large scale MOS integrated circuits E. Lelarasmee, A. Sangiovanni-Vincentelli

January 1982 Proceedings of the 19th conference on Design automation

Full text available: pdf(755.47 KB)

Additional Information: full citation, abstract, references, citings, index terms

Algorithms and techniques used in RELAX are described. RELAX is a time domain MOS digital circuit simulator based on a new analysis method called Waveform Relaxation Method [1] which exploits decomposition techniques. Preliminary comparisons between RELAX and the standard circuit simulator SPICE2 have shown that RELAX is a fast and reliable circuit simulator.

18 Implementation of sparta, a highly parallel circuit simulator by the preconditioned Jacobi method, on a distributed memory machine Reiji Suda, Yoshio Oyanagi

July 1995 Proceedings of the 9th international conference on Supercomputing

Full text available: pdf(825.77 KB) Additional Information: full citation, references, index terms

19 RELAX: A new circuit simulator for large scale MOS integrated circuits E. Lelarasmee, A. Sangiovanni-Vincentelli June 1988 Papers on Twenty-five years of electronic design automation

Full text available: pdf(890.06 KB) Additional Information: full citation, references, index terms

²⁰ Survey of analysis, simulation and modeling for large scale logic circuits Albert E. Ruehli

June 1981 Proceedings of the 18th conference on Design automation

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(481.15 KB) terms

The purpose of this paper is to introduce recent developments in the time analysis, simulation and modeling of logic circuits. These advances which have taken place in the circuit and systems area augment the recent advances in logic time simulators. The latest trend has been to combine the approaches into a single system, a so called mixed simulation-analysis program. In this paper we review some of the circuit oriented techniques at a level understandable to the non circuit-theorist.

Results 1 - 20 of 200 Result page: 1 2 3 4 5 6 7 8 9 10

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2004 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Mindows Media Player Real Player



Search: • The ACM Digital Library • The Guide

+cas +latency

1201590

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survev

Terms used cas latency

Found 474 of 145,831

Sort results

by

Display results

relevance

expanded form

Save results to a Binder Search Tips

Open results in a new window

Try an Advanced Search Try this search in The ACM Guide

Results 21 - 40 of 200

Result page: previous 1 **2** 3 4 5 6 7 8 9 10

next Relevance scale

Best 200 shown

21 Low-energy off-chip SDRAM memory systems for embedded applications Hojun Shim, Yongsoo Joo, Yongseok Choi, Hyung Gyu Lee, Naehyuck Chang February 2003 ACM Transactions on Embedded Computing Systems (TECS), Volume 2 Issue

Full text available: pdf(3.98 MB)

Additional Information: full citation, abstract, references, index terms

Memory systems are dominant energy consumers, and thus many energy reduction techniques for memory buses and devices have been proposed. For practical energy reduction practices, we have to take into account the interaction between a processor and cache memories together with application programs. Furthermore, energy characterization of memory systems must be accurate enough to justify various techniques. In this article, we build an in-house energy simulator for memory systems that is accelerat ...

Keywords: Low power, SDRAM, memory system

22 Design space exploration for embedded systems: Energy exploration and reduction of SDRAM memory systems



Yongsoo Joo, Yongseok Choi, Hojun Shim, Hyung Gyu Lee, Kwanho Kim, Naehyuck Chang June 2002 Proceedings of the 39th conference on Design automation

Full text available: pdf(196.08 KB)

Additional Information: full citation, abstract, references, citings, index terms

In this paper, we introduce a precise energy characterization of SDRAM main memory systems and explore the amount of energy associated with design parameters, leading to energy reduction techniques that we are able to recommend for practical use. We build an in-house energy simulator for SDRAM main memory systems based on cycle-accurate energy measurement and state-machine-based characterizations which independently characterize dynamic and static energy. We explore energy behavior of the memory ...

Keywords: SDRAM, low power, memory system

23 Leakage Power Optimization Techniques for Ultra Deep Sub-Micron Multi-Level



Nam Sung Kim, David Blaauw, Trevor Mudge

November 2003 Proceedings of the 2003 international conference on on Computeraided design

Full text available: pdf(182.94 KB) Additional Information: full citation, abstract

On-chip L1 and L2 caches represent a sizeable fraction of the totalpower consumption of microprocessors. In deep sub-micron technology, the subthreshold leakage power is becoming the dominantfraction of the total power consumption of those caches. In thispaper, we present optimization techniques to reduce the leakagepower of on-chip caches assuming that there are multiple thresholdvoltages, VTH's, available. First, we show a cache leakage optimizationtechnique that examines the trade-off between ...

24 Synchronization mechanisms for SCRAMNet+ systems

Stephen Menke, Mark Moir, Srikanth Ramamurthy

June 1998 Proceedings of the seventeenth annual ACM symposium on Principles of distributed computing

Full text available: pdf(1.35 MB)

Additional Information: full citation, references, index terms

25 Technicial session 5: student best paper contest: Predictive perceptual compression



Oleg Komogortsev, Javed Khan

October 2004 Proceedings of the 12th annual ACM international conference on Multimedia

Full text available: pdf(514.78 KB) Additional Information: full citation, abstract, references, index terms

Approximately 2 degrees in our 140 degree vision span has sharp vision. Many researchers have been fascinated by the idea of eye-tracking integrated perceptual compression of an image or video, yet any practical system has yet to emerge. The unique challenge presented by real time perceptual video streaming is how to handle the fast nature of the human eye and provide its integration with computationally intensive video transcoding scheme. The delay introduced by video transmission in the net ...

Keywords: perceptual compression, video transcoding

26 Dynamically allocating processor resources between nearby and distant ILP Rajeev Balasubramonian, Sandhya Dwarkadas, David H. Albonesi

May 2001 ACM SIGARCH Computer Architecture News, Proceedings of the 28th annual international symposium on Computer architecture, Volume 29 Issue 2

Full text available: pdf(998.02 KB)

Additional Information: full citation, abstract, references, citings, index terms

Modern superscalar processors use wide instruction issue widths and out-of-order execution in order to increase instruction-level parallelism (ILP). Because instructions must be committed in order so as to quarantee precise exceptions, increasing ILP implies increasing the sizes of structures such as the register file, issue queue, and reorder buffer. Simultaneously, cycle time constraints limit the sizes of these structures, resulting in conflicting design requirements.

In ...

²⁷ Reliable communications in FTL

Ivan Kalas

November 1995 Proceedings of the 1995 conference of the Centre for Advanced Studies on Collaborative research

Full text available: pdf(245.46 KB) Additional Information: full citation, abstract, references, index terms

Local-area networks based on high-bandwidth packet-switching technology, such as ATM,



show a tremendous promise in reducing communication latencies and overheads. However, the lack of flow-control and reliable delivery in ATM networks requires the higher protocol layers to deal with cell loss or corruption. While it is possible to use TCP-based communications over ATM, the protocol mismatch results in a significant loss of the bandwidth. In addition to this, there is also a significant mismatch ...

28 Efficient synchronization for nonuniform communication architectures

Zoran Radović, Erik Hagersten

November 2002 Proceedings of the 2002 ACM/IEEE conference on Supercomputing

Full text available: 📆 pdf(162.38 KB) Additional Information: full citation, abstract, references, index terms

Scalable parallel computers are often nonuniform communication architectures (NUCAs), where the access time to other processor's caches vary with their physical location. Still, few attempts of exploring cache-to-cache communication locality have been made. This paper introduces a new kind of synchronization primitives (lock-unlock) that favor neighboring processors when a lock is released. This improves the lock handover time as well as access time to the shared data of the critical region. A cr ...

29 Shared memory and architecture: A scalable lock-free stack algorithm

Danny Hendler, Nir Shavit, Lena Yerushalmi

June 2004 Proceedings of the sixteenth annual ACM symposium on Parallelism in algorithms and architectures

Full text available: pdf(221.87 KB) Additional Information: full citation, abstract, references, index terms

The literature describes two high performance concurrent stack algorithms based on combining funnels and elimination trees. Unfortunately, the funnels are linearizable but blocking, and the elimination trees are non-blocking but not linearizable. Neither is used in practice since they perform well only at exceptionally high loads. The literature also describes a simple lock-free linearizable stack algorithm that works at low loads but does not scale as the load increases. The question of designi ...

30 A performance comparison of contemporary DRAM architectures

Vinodh Cuppu, Bruce Jacob, Brian Davis, Trevor Mudge

May 1999 ACM SIGARCH Computer Architecture News, Proceedings of the 26th annual international symposium on Computer architecture, Volume 27 Issue 2

Full text available: pdf(166.88 KB) Additional Information: full citation, abstract, references, citings, index terms Publisher Site

In response to the growing gap between memory access time and processor speed, DRAM manufacturers have created several new DRAM architectures. This paper presents a simulation-based performance study of a representative group, each evaluated in a small system organization. These small-system organizations correspond to workstation-class computers and use on the order of 10 DRAM chips. The study covers Fast Page Mode, Extended Data Out, Synchronous, Enhanced Synchronous, Synchronous Link, Rambus, ...

31 Network behavior: An analysis of Internet content delivery systems Stefan Saroiu, Krishna P. Gummadi, Richard J. Dunn, Steven D. Gribble, Henry M. Levy December 2002 ACM SIGOPS Operating Systems Review, Volume 36 Issue SI

Full text available: pdf(2.07 MB) Additional Information: full citation, abstract, references

In the span of only a few years, the Internet has experienced an astronomical increase in the use of specialized content delivery systems, such as content delivery networks and peerto-peer file sharing systems. Therefore, an understanding of content delivery on the Internet now requires a detailed understanding of how these systems are used in practice. This paper examines content delivery from the point of view of four content delivery systems: HTTP web traffic, the Akamai content delivery netw ...

32 An empirical analysis of the performance of a multiprocessor-based circuit simulator George K. Jacob, A. Richard Newton, Donald O. Pederson July 1986 Proceedings of the 23rd ACM/IEEE conference on Design automation



Full text available: pdf(727.00 KB)

Additional Information: full citation, abstract, references, citings, index terms

Our original MSPLICE multiprocessor-based circuit simulator showed excellent efficiency with up to 10 processors. As shown in this paper, however, the efficiency of the program drops significantly when over 40 processors are used. A new generation of the MSPLICE program is described which shows high efficiency with up to 99 processors for three different benchmark circuits. Data is compared against predictions made from simulations of an ideal Gauss-Seidel machine model with unit delay, and ...

33 Robust interfaces for mixed-timing systems with application to latency-insensitive protocols



Tiberiu Chelcea, Steven M. Nowick

June 2001 Proceedings of the 38th conference on Design automation

Full text available: pdf(124.41 KB)

Additional Information: full citation, abstract, references, citings, index <u>terms</u>

This paper presents several low-latency mixed-timing FIFO designs that interface systems on a chip working at different speeds. The connected systems can be either synchronous or asynchronous. The design are then adapted to work between systems with very long interconnection delays, by migrating a single-clock solution by Carloni et al. (for "latencyinsensitive" protocols) to mixed-timing domains. The new designs can be made arbitrarily robust with regard to metastability and i ...

34 VHDL analog extensions: process, issues and status

Robert Cottrell, Kevin Nolan, Mark Brown

November 1992 Proceedings of the conference on European design automation

Full text available: pdf(582.83 KB) Additional Information: full citation, references, index terms

35 Concurrency, latency, or system overhead: which has the largest impact on uniprocessor DRAM-system performance?

Vinodh Cuppu, Bruce Jacob

May 2001 ACM SIGARCH Computer Architecture News, Proceedings of the 28th annual international symposium on Computer architecture, Volume 29 Issue 2

Full text available: pdf(904.17 KB)

Additional Information: full citation, abstract, references, citings, index <u>terms</u>

Given a fixed CPU architecture and a fixed DRAM timing specification, there is still a large design space for a DRAM system organization. Parameters include the number of memory channels, the bandwidth of each channel, burst sizes, queue sizes and organizations, turnaround overhead, memory-controller page protocol, algorithms for assigning request priorities and scheduling requests dynamically, etc. In this design space, we see a wide variation in application execution times: for example, ...

36 Exploiting cache affinity in software cache coherence

Hui Li, Kenneth C. Sevcik

July 1994 Proceedings of the 8th international conference on Supercomputing

Full text available: pdf(999.44 KB) Additional Information: full citation, abstract, references, index terms

Cache affinity is important to the performance of scalable shared memory multiprocessors.

For multiprocessors without hardware cache coherence support, software cache coherence is the only alternative. Most existing software cache schemes ignore cache affinity across parallel loops. In this paper, we propose a new scheme, Cache Affinity-based Software cache coherence scheme (CAS), that exploits cache affinity across parallel loops to achieve high cache hit ratios without requiring extra har ...

Latency and latch count minimization in wave steered circuits
 Amit Singh, Arindam Mukherjee, Malgorzata Marek-Sadowska
 June 2001 Proceedings of the 38th conference on Design automation

Full text available: pdf(151.69 KB)

Additional Information: full citation, abstract, references, citings, index terms

Wave Steering is a new design methodology that realizes high throughput circuits by embedding layout friendly synthesized structures in silicon. Wave Steered circuits inherently utilize latches in order to guarantee the correct signal arrival times at the inputs of these synthesized structures and maintain the high throughput of operation. In this paper, we show a method of reor-dering signals to achieve minimum circuit latency for Wave Steered circuits and propose an Integer Linear Program ...

38 Performance analysis and optimization of latency insensitive systems Luca P. Carloni, Alberto L. Sangiovanni-Vincentelli June 2000 Proceedings of the 37th conference on Design automation

Full text available: pdf(235.41 KB)

Additional Information: full citation, abstract, references, citings, index terms

Latency insensitive design has been recently proposed in literature as a way to design complex digital systems, whose functional behavior is robust with respect to arbitrary variations in interconnect latency. However, this approach does not guarantee the same robustness for the performance of the design, which indeed can experience big losses. This paper presents a simple, yet rigorous, method to (1) model the key properties of a latency insensitive system, (2) analyze the impact o ...

39 The Mercury Interconnect Architecture: a cost-effective infrastructure for highperformance servers

Wolf-Dietrich Weber, Stephen Gold, Pat Helland, Takeshi Shimizu, Thomas Wicki, Winfried Wilcke

May 1997 ACM SIGARCH Computer Architecture News, Proceedings of the 24th annual international symposium on Computer architecture, Volume 25 Issue 2

Full text available: pdf(1.53 MB)

Additional Information: full citation, abstract, references, citings, index terms

This paper presents HAL's Mercury Interconnect Architecture, an interconnect infrastructure designed to link commodity microprocessors, memory, and I/O components into high-performance multiprocessing servers. Both shared-memory and message-passing systems, as well as hybrid systems are supported by the interconnect. The key attributes of the Mercury Interconnect Architecture are: low latency, high bandwidth, a modular and flexible design, reliability/availability/serviceability (RAS) features, ...

40 Architectural retiming: pipelining latency-constrained circuits
Soha Hassoun, Carl Ebeling

June 1996 Proceedings of the 33rd annual conference on Design automation

Full text available: pdf(404.79 KB) Additional Information: full citation, references, citings, index terms

Results 21 - 40 of 200 Result page: <u>previous</u> <u>1</u> <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u> <u>next</u>